Appendix A: Basic Notions of Ergodic Theory

In this chapter we summarize some of the background in ergodic theory, referring to [53] for a thorough treatment. We also describe a few examples that will be studied frequently in this volume.

A.1 Basic Properties of Measure-Preserving Systems

The primary object in ergodic theory is a *measure-preserving system*, consisting of a probability space $(X, \mathscr{B}, \mu)$ together with a measure-preserving transformation $T : X \to X$. This means that:

- X is a set, whose elements will be referred to as *points*;
- $\mathscr{B}$ is a σ-algebra of subsets of X, whose elements are the *measurable sets*;
- $T^{-1}(B) \in \mathscr{B}$ for all $B \in \mathscr{B}$; and
- μ is a probability measure defined on $\mathscr{B}$ which is *preserved* by T in the sense that $\mu(T^{-1}B) = \mu(B)$ for all $B \in \mathscr{B}$.

We also say the measure μ is invariant for T. A fundamental property of any such system is *Poincaré recurrence*, which we phrase in two equivalent ways.

(1) For any set $A \in \mathscr{B}$ with $\mu(A) > 0$, there is some n with $1 \leqslant n \leqslant \frac{1}{\mu(A)}$ for which $\mu(A \cap T^{-n}A) > 0$. Colloquially, every set of positive measure returns to itself non-trivially.
(2) For any set $E \in \mathscr{B}$ there is a measurable set $F \subseteq E$ with $\mu(F) = \mu(E)$ with the property that for every $x \in F$ there exist integers $0 < n_1 < n_2 < \cdots$ with $T^{n_i}x \in E$ for all $i \geqslant 1$. Colloquially, almost every point of any measurable set returns to the set infinitely often.

Poincaré recurrence is a form of pigeon-hole principle in ergodic theory, and if the measure space is isomorphic to a space with finitely many points, then it is the pigeon-hole principle. It lies at the bottom of a hierarchy of recurrence phenomena.

Two measure-preserving systems may relate to each other in many ways, the simplest of which is that they may be isomorphic from the measure-theoretic point of view, or that one may be an image of the other from the measure-theoretic point of view. Given $(X, \mathscr{B}_X, \mu, T)$ and $(Y, \mathscr{B}_Y, \nu, S)$, two measure-preserving systems, we make the following definitions.

(1) The system $(Y, \mathscr{B}_Y, \nu, S)$ is a *factor* of $(X, \mathscr{B}_X, \mu, T)$ if there are sets X' in $\mathscr{B}_X$ and Y' in $\mathscr{B}_Y$ with $\mu(X') = 1$, $\nu(Y') = 1$, $TX' \subseteq X'$, $SY' \subseteq Y'$ and a measure-preserving map $\phi : X' \to Y'$ with

$$\phi \circ T(x) = S \circ \phi(x)$$

for all $x \in X'$.

(2) The system $(Y, \mathscr{B}_Y, \nu, S)$ is *isomorphic to* $(X, \mathscr{B}_X, \mu, T)$ if there are sets X' in $\mathscr{B}_X$, Y' in $\mathscr{B}_Y$ with $\mu(X')=1$, $\nu(Y')=1$, $TX' \subseteq X'$, $SY' \subseteq Y'$, and an invertible[†] measure-preserving map $\phi : X' \to Y'$ with

$$\phi \circ T(x) = S \circ \phi(x)$$

for all $x \in X'$.

It is natural to simply ignore null sets (because they are not visible from a measure-theoretic point of view), and so we may sometimes think of a factor as a measure-preserving map $\phi : X \to Y$ for which the diagram

$$\begin{array}{ccc} X & \xrightarrow{\;T\;} & X \\ {\scriptstyle\phi}\big\downarrow & & \big\downarrow{\scriptstyle\phi} \\ Y & \xrightarrow[\;S\;]{} & Y \end{array}$$

is commutative, with the understanding that the map is not required to be defined everywhere.

A factor map

$$(X, \mathscr{B}_X, \mu, T) \longrightarrow (Y, \mathscr{B}_Y, \nu, S)$$

is also called an *extension* of $(Y, \mathscr{B}_Y, \nu, S)$. The factor $(Y, \mathscr{B}_Y, \nu, S)$ is called *trivial* if as a measure space Y comprises a single element; the extension is called *trivial* if ϕ is an isomorphism of measure spaces.

One of the main internal problems in ergodic theory is the classification of measure-preserving systems up to the equivalence relation of measurable isomorphism. While this has been achieved for special kinds of measure-preserving systems, at this level of generality no reasonable description can be given (we refer to Section 1.7 for a brief discussion and references). The study

[†] This means that we require the map to be measurable, to have an inverse, and that the inverse is measurable.

of internal questions of this sort is a vibrant part of ergodic theory nonetheless, and we refer the interested reader to the monographs of Rudolph [178] and Glasner [68] for more on this.

Our primary interests lie elsewhere, for example in the connections between ergodic theory and number theory. Many of these connections arise through the study of specific examples of measure-preserving systems.

A.1.1 The Associated Operator on L^2_μ

Any measure-preserving map T on a probability space $(X, \mathscr{B}, \mu)$ has an associated operator $U_T : L^2_\mu \to L^2_\mu$ defined by

$$U_T(f) = f \circ T.$$

Notice that

$$\begin{aligned}\langle U_T f_1, U_T f_2 \rangle &= \int f_1 \circ T \cdot \overline{f_2 \circ T}\,\mathrm{d}\mu \\ &= \int f_1 \overline{f_2}\,\mathrm{d}\mu \qquad \text{(since } \mu \text{ is } T\text{-invariant)} \\ &= \langle f_1, f_2 \rangle,\end{aligned}$$

so U_T is an isometry mapping L^2_μ into L^2_μ (here $\langle \cdot, \cdot \rangle$ is the usual inner-product on L^2_μ). If T is invertible (so that T^{-1} is also a measure-preserving map), then $U_T \circ U_{T^{-1}} = U_{T^{-1}} \circ U_T$ is the identity map on L^2_μ, and so the associated operator U_T is unitary†. A property of a measure-preserving system is called *spectral* or *unitary* if it can be detected by studying the associated operator on L^2_μ alone.

A.1.2 Ergodic Theorems

The origins of ergodic theory lie partly in physical questions, and the ergodic theorems find some of their motivations there. The conclusions of the ergodic theorems, two of which we will state here, are properties of any measure-preserving system, and in particular are not connected to the hypothesis of ergodicity described below. Under the assumption of ergodicity, the conclusions are strengthened. The first is the mean ergodic theorem, due to von Neumann [142]. It predates the pointwise ergodic theorem of Birkhoff [13].

Theorem A.1 (Mean Ergodic Theorem). *Let $(X, \mathscr{B}, \mu, T)$ be a measure-preserving system, and let P_T denote the orthogonal projection onto the closed subspace*

$$I = \{g \in L^2_\mu \mid U_T g = g\} \subseteq L^2_\mu.$$

Then for any $f \in L^2_\mu$,

† That is, the adjoint satisfies $U_T^* = U_T^{-1}$.

$$\frac{1}{N}\sum_{n=0}^{N-1} U_T^n f \underset{L^2_\mu}{\longrightarrow} P_T f.$$

The mean ergodic theorem does not say anything about the averages along individual orbits of the value of the function f (indeed f itself is an equivalence class of functions, so it does not make sense to speak about its value at any given point; such expressions are often called *ergodic averages*). The next result, the pointwise or individual ergodic theorem, describes the behavior of ergodic averages for almost every point. In order to state it, we introduce the notation $\mathscr{L}^1_\mu$ for the space of integrable measurable functions on $(X, \mathscr{B}, \mu)$ (as opposed to the space L^1_μ of equivalence classes of functions).

Theorem A.2 (Pointwise ergodic theorem). *Let $(X, \mathscr{B}, \mu, T)$ be a measure-preserving system. If $f \in \mathscr{L}^1_\mu$, then the ergodic averages*

$$\frac{1}{n}\sum_{j=0}^{n-1} f(T^j x)$$

converge almost everywhere and in L^1_μ as $n \to \infty$ to a T-invariant function $f^* \in \mathscr{L}^1_\mu$, and

$$\int f^* \,\mathrm{d}\mu = \int f \,\mathrm{d}\mu.$$

A.1.3 Basic Properties

Just as Poincaré recurrence is the first of many recurrence properties, ergodicity is the simplest of many so-called mixing properties. A measure-preserving system $(X, \mathscr{B}, \mu, T)$ is *ergodic* if any of the following equivalent properties hold.

(1) For any $B \in \mathscr{B}$, $T^{-1}B = B \implies \mu(B) = 0$ or $\mu(B) = 1$.
(2) For any $B \in \mathscr{B}$, $\mu\left(B \triangle T^{-1}B\right) = 0 \implies \mu(B) = 0$ or $\mu(B) = 1$.
(3) For $A \in \mathscr{B}$, $\mu(A) > 0$ implies that $\mu\left(\bigcup_{n=1}^{\infty} T^{-n}A\right) = 1$.
(4) For $A, B \in \mathscr{B}$, $\mu(A)\mu(B) > 0$ implies that there exists $n \geqslant 1$ with

$$\mu(T^{-n}A \cap B) > 0.$$

(5) For $f : X \to \mathbb{C}$ measurable, $f \circ T = f$ almost everywhere implies that f is equal to a constant almost everywhere.
(6) For $f \in L^2_\mu$, $U_T f = f$ as elements of L^2_μ implies that f is equal to a constant in L^2_μ.

The hypothesis on the set B in (1) is called *strict invariance*, while (2) is usually called *invariance* under T (see Exercise A.1.1). The hypothesis in (2) will be useful: we write $A = B \pmod{\mu}$ or

$$A \underset{\mu}{=} B$$

to mean that $\mu(A \triangle B) = 0$.

The limits in the mean and pointwise ergodic theorems are invariant functions, so both theorems can be strengthened under the hypothesis of ergodicity as follows.

(1) If $(X, \mathscr{B}, \mu, T)$ is ergodic then

$$P_T f = \int f \, \mathrm{d}\mu$$

as elements of L^2_μ.

(2) If $(X, \mathscr{B}, \mu, T)$ is ergodic then

$$f^*(x) = \int f \, \mathrm{d}\mu$$

for almost every $x \in X$ with respect to μ.

Ergodicity ensures that any two sets A and B of positive measure *mix* under the action of T in the weakest possible sense, namely that there is some time $n \geqslant 1$ for which A and $T^{-n}B$ intersect non-trivially. Requiring this to happen in a more controlled way provides two further notions. A measure-preserving system $(X, \mathscr{B}, \mu, T)$ is said to be *weak-mixing* if

$$\frac{1}{N} \sum_{n=0}^{N-1} \left| \mu(A \cap T^{-n}B) - \mu(A)\mu(B) \right| \longrightarrow 0$$

as $N \to \infty$, for all $A, B \in \mathscr{B}$. Weak-mixing can be phrased in a surprising number of seemingly different but equivalent ways (see [53, Sect. 2.7–8] for the details). Of these equivalent formulations, we only note here that T is weak-mixing if and only if the associated operator U_T has no measurable eigenfunctions other than the constant functions (that is, T has continuous spectrum). In particular, weak-mixing is a unitary property.

A stronger property is to require that sets become asymptotically independent under the action of T, and $(X, \mathscr{B}, \mu, T)$ is said to be *mixing* if

$$\mu\left(A \cap T^{-n}B\right) \to \mu(A)\mu(B)$$

as $n \to \infty$ for any sets $A, B \in \mathscr{B}$. Once again there are equivalent formulations of this notion, but mixing has nothing like the range of equivalent formulations of weak-mixing. The following properties are equivalent for a measure-preserving system $(X, \mathscr{B}, \mu, T)$.

(1) $(X, \mathscr{B}, \mu, T)$ is mixing.
(2) $\langle U_T^n f_1, f_2 \rangle \to \langle f_1, 1 \rangle \cdot \langle 1, f_2 \rangle$ as $n \to \infty$ for any $f_1, f_2 \in L^2_\mu$.

(3) $\langle U_T^n f, f\rangle \to 0$ as $n \to \infty$ for all f in a set of functions dense in the set of all functions in L^2_μ with zero integral.

All three notions are distinct: there are measure-preserving systems that are ergodic but not weak-mixing, and there are systems that are weak-mixing but not mixing.

Exercises for Section A.1

Exercise A.1.1. Let $B \in \mathscr{B}$ be an invariant set in a measure-preserving system $(X, \mathscr{B}, \mu, T)$. Show that there is a strictly invariant set $C \in \mathscr{B}$ with $\mu(C \triangle B) = 0$.

Exercise A.1.2. Prove Kac's quantitative form of Poincaré recurrence theorem [92] in the following form. Let $(X, \mathscr{B}, \mu, T)$ be an invertible ergodic measure-preserving system, and let $A \in \mathscr{B}$ have $\mu(A) > 0$.
(a) Show that

$$\int_A r_A(x)\,\mathrm{d}\mu(x) = 1 \tag{A.1}$$

where r_A is the return time function defined almost everywhere on A by

$$r_A(x) = \inf\{m \geqslant 1 \mid T^m x \in A\}.$$

Equivalently, the expected return time to A is $\frac{1}{\mu(A)}$.
(b) Show that the *first return map* $T_A : A \to A$ defined almost everywhere by

$$T_A(x) = T^{r_A(x)}(x)$$

is measure-preserving and ergodic with respect to the normalized restricted measure $\mu_A = \frac{1}{\mu(A)}\mu|_A$.
(c) What form of (A.1) holds without the assumption of ergodicity?
(d) Show that if an invertible measure-preserving system satisfies (A.1) for all sets A of positive measure, then it is ergodic.

Exercise A.1.3. (Quantitative almost sure recurrence[(47)]) Let (X, d) be a compact metric space with finite upper box dimension†. Assume that

$$T : X \to X$$

† This means that there exists some $d > 0$ such that for every $\varepsilon > 0$ and every sufficiently small $r > 0$ there exist points $x_1, \dots, x_{J(r)} \in X$ with $J(r) \leqslant e^{-(d+\varepsilon)}$ and with

$$X = \bigcup_{j=1}^{J(r)} B_r(x_j).$$

is a continuous map preserving a Borel probability measure μ on X. Show that for μ-almost every $x \in X$ and all $\varepsilon > 0$ we have

$$\min\{n \geqslant 1 \mid \mathsf{d}(x, T^n x) < r\} \leqslant r^{-(d+\varepsilon)}$$

for all sufficiently small $r > 0$.

Exercise A.1.4. (48) Let $(X, \mathscr{B}, \mu, T)$ be a measure-preserving system, with X a compact metric space and $\mathscr{B}$ the Borel σ-algebra on X, satisfying the two properties

- T^k is ergodic for all $k \geqslant 1$; and
- there is a constant $C > 1$ with

$$\limsup_{n\to\infty} \mu\left(T^{-n}A \cap B\right) \leqslant C\mu(A)\mu(B)$$

  for all sets A and B in $\mathscr{B}$.

Show that the system is mixing. Also show that it is sufficient to have the second property for all sets A, B in a semi-algebra that generates $\mathscr{B}$.

A.2 Continuous Maps on Metric Spaces

For the material reviewed here, we refer to [53, Ch. 4 and Sect. B.5] for a detailed treatment in the compact case†.

A natural family of dynamical systems are given by continuous maps

$$T : X \to X,$$

where $X = (X, \mathsf{d})$ is a compact metric space. We write $\mathscr{M}(X)$ for the space of Borel probability measures on X, which is a compact convex metrizable space in the weak*-topology. The map T induces a continuous map

$$T_* : \mathscr{M}(X) \to \mathscr{M}(X)$$

by defining the image $T_*\mu$ of a measure μ by $T_*\mu(A) = \mu(T^{-1}A)$. Write $\mathscr{M}^T(X)$ for the subset of $\mathscr{M}(X)$ comprising the measures invariant under T, and $\mathscr{E}^T(X)$ for the subset of ergodic measures. The main properties of this space are as follows.

(1) The space $\mathscr{M}^T(X)$ is a non-empty(49) compact convex set.
(2) The ergodic invariant measures are exactly the extreme points of $\mathscr{M}^T(X)$.
(3) If $\mu, \nu \in \mathscr{E}^T(X)$ then either $\mu = \nu$ or μ and ν are mutually singular.

† We will mostly be concerned with the compact case, but for completeness we refer to Stroock [192] or Bourbaki [20] for the general case.

Because of the first property, any continuous map of a compact metric space generates one, and possibly many, measure-preserving systems. The identification of measures with prescribed dynamical properties is a recurrent theme in this volume.

Many of the important examples (notably, some translations on quotients of Lie groups, and interval exchange transformations) are not compact. If X is a locally compact space, then we write $\mathscr{M}(X)$ for the space of Borel probability measures on X. This space may be non-compact in the weak topology. For a continuous map $T : X \to X$, we again write $\mathscr{M}^T(X)$ for the space of T-invariant measures, and this space may be empty. Also complications may arise when the space is not empty. For instance in this case one can study the orbit of a T-invariant measure under some other transformation, and ask whether there is escape of mass (that is, whether limit points of the orbit are still probability measures).

The ergodic decomposition shows that every invariant measure can be decomposed into a generalized convex combination of ergodic invariant measures — and so ergodic measures are omnipresent. As the ergodic measures are precisely the extremal points of the convex and compact set of invariant measures, the existence of an ergodic decomposition can be derived from Choquet's theorem. We prefer the alternate approach using conditional measures as these also play a central role in the study of entropy, see Chapter 2 and in particular Section 2.1.4.

In contrast to the ergodic decomposition, there is no sense in which an arbitrary measure-preserving system can be decomposed into mixing or weak-mixing components.

Exercises for Section A.2

Exercise A.2.1. Let $T : X \to X$ be a continuous map on a compact metric space, and assume that (ν_n) is a sequence of T-invariant Borel probability measures. Suppose that

$$\frac{1}{N}\sum_{n=1}^{N}\nu_n \longrightarrow \mu$$

as $N \to \infty$ in the weak*-topology, where μ is an ergodic measure.
(a) Show that there is a subset $S \subseteq \mathbb{N}$ with density one with the property that

$$\nu_n \underset{n\in S}{\longrightarrow} \mu$$

in the weak*-topology.
(b) Show that the same conclusion holds without the assumption that each measure ν_n is T-invariant but with the weaker assumption that

$$\left|\int f\circ T\,\mathrm{d}\nu_n - \int f\,\mathrm{d}\nu_n\right| \longrightarrow 0$$

as $n \to \infty$ for all continuous functions $f \in C(X)$ (this weaker hypothesis is sometimes called *asymptotic invariance*).

Exercise A.2.2. Let $T : X \to X$ and $S : Y \to Y$ be continuous maps of compact metric spaces, and let $\phi : X \to Y$ be a topological factor map (as in Definition 5.27).
(a) If ν is an S-invariant Borel probability measure on Y, show that there exists a T-invariant Borel probability measure μ on X such that $\phi_*\mu = \nu$.
(b) If in addition ν is ergodic with respect to S, show that μ can be chosen to be ergodic as well.

A.3 Examples

As mentioned above, examples play a significant role in ergodic theory and dynamical systems for several reasons. One reason is that there are relatively few general statements about measure-preserving systems or continuous maps with no additional hypotheses(50). A second reason is that many of the most powerful applications of ergodic theory, or of ergodic ideas, arise through specific examples. In this short section we assemble a small collection of examples, several of which are discussed in much greater detail in [53], and some of which will feature heavily in the remainder of this book.

Example A.3. Let G be a compact metrizable abelian group with Haar measure m_G, and fix $g \in G$. Then the map $R_g : G \to G$ defined by $R_g(h) = hg$ is called the *group rotation* by g.

- The map R_g is continuous, and m_G is invariant under R_g.
- The measure m_G is ergodic if and only if the smallest closed subgroup of G containing g is G itself.
- If R_g is ergodic, then it is uniquely ergodic in the sense that $|\mathscr{E}^T(G)| = 1$.
- The measure-preserving system (G, R_g, μ) is not weak-mixing.

A special case of Example A.3 is a circle rotation $t \mapsto t+\theta$ modulo 1 on the circle $\mathbb{R}/\mathbb{Z}$ identified with $[0, 1)$ under addition modulo 1. This map preserves Lebesgue measure (Haar measure for the circle), and Lebesgue measure is ergodic if and only if $\theta \notin \mathbb{Q}$.

Example A.4. Let G be a compact metrizable abelian group with Haar measure m_G, and let T be a continuous endomorphism (surjective homomorphism) from G to G.

- The measure m_G is invariant under T.
- The measure m_G is ergodic for T if and only if the identity $\chi(T^n x) = \chi(x)$ for some $n \geqslant 1$ and character $\chi \in \widehat{X}$ implies that χ is the trivial character.
- If $|G| > 1$ then the map T is not uniquely ergodic, since both Haar measure m_G and the point mass at the identity δ_{I_G} are invariant.

- If T is ergodic with respect to m_G, then T is mixing with respect to m_G.

Two special cases of Example A.4 are particularly important as model systems.

(1) If $G = \mathbb{R}/\mathbb{Z} = \mathbb{T}$, the circle, and $T(x) = 2x$ modulo 1, then T is called the circle doubling map. It is ergodic and hence mixing.
(2) Let $G = \mathbb{R}^d/\mathbb{Z}^d = \mathbb{T}^d$, the d-dimensional torus, and let $A \in \mathrm{Mat}_{dd}(\mathbb{Z})$ be an integer matrix with $\det(A) \neq 0$. Then A induces a surjective homomorphism $T_A : \mathbb{T}^d \to \mathbb{T}^d$ which preserves Lebesgue measure. The map is ergodic with respect to the Lebesgue measure if and only if no eigenvalue of A is a root of unity. If $A \in \mathrm{GL}_d(\mathbb{Z})$ then T_A is a *toral automorphism*.

The next collection of examples arise whenever a dynamical system is studied via a partition. We will follow the treatment of Lind and Marcus [119] closely, and refer to their book for a complete treatment.

Example A.5. Let A be a finite set called an *alphabet*, and let F be a collection of blocks (finite concatenations of symbols from the alphabet A). The space $A^{\mathbb{Z}}$ is compact in the product topology (the alphabet A is given the discrete topology). A convenient metric yielding this topology is given by

$$\mathsf{d}(x, y) = \sum_{n \in \mathbb{Z}} |x_i - y_i| \cdot 2^{-|n|}. \tag{A.2}$$

Thus two points are close in this metric if their coordinates agree on a large neighborhood of the position 0. The sets of the form

$$[w]_m^n = \{x \in A^{\mathbb{Z}} \mid x_i = w_i \text{ for } i = m, \dots, n\}$$

are called *cylinder sets* and for $m = -n$ these are the metric open balls. For convenience we will write $[w]_n$ for $[w]_n^n$. Notice that in the definition of a cylinder set w could be either a finite block of $n - m + 1$ symbols from the alphabet, or an element of $A^{\mathbb{Z}}$ itself.

Define the subset X_F of $A^{\mathbb{Z}}$ to be the set of all bi-infinite sequences of symbols from A with the property that no block from F appears anywhere in the sequence. A set of this form is called a *shift space*, and it is closed (and hence compact) in the metric topology. The cylinder sets are defined in a shift space by intersecting a cylinder set in the full shift with X_F. The left shift $\sigma : \mathsf{X}_\mathsf{F} \to \mathsf{X}_\mathsf{F}$ defined by

$$(\sigma x)_k = x_{k+1}$$

for all $k \in \mathbb{Z}$, is a homeomorphism of X_F. The map $\sigma : \mathsf{X}_\mathsf{F} \to \mathsf{X}_\mathsf{F}$ is called a *shift dynamical system*.

A shift-commuting map between a shift $\mathsf{X}_{\mathsf{F}_1}$ with shift map σ_1 over an alphabet A_1 and a shift $\mathsf{X}_{\mathsf{F}_2}$ with shift map σ_2 over an alphabet A_2 is a map $\phi : \mathsf{X}_{\mathsf{F}_1} \to \mathsf{X}_{\mathsf{F}_2}$ with the property that the diagram

$$\begin{array}{ccc} \mathsf{X}_{\mathsf{F}_1} & \xrightarrow{\sigma_1} & \mathsf{X}_{\mathsf{F}_1} \\ \phi\big\downarrow & & \big\downarrow\phi \\ \mathsf{X}_{\mathsf{F}_2} & \xrightarrow[\sigma_2]{} & \mathsf{X}_{\mathsf{F}_2} \end{array}$$

commutes. Notice that a shift-commuting map ϕ is determined by its zero coordinate $\psi(x) = (\phi(x))_0 \in A_2$, since

$$(\phi(x))_k = \psi(\sigma_1^k x)$$

for all $k \in \mathbb{Z}$ (that is, $\phi(x)$ is the element of $\mathsf{X}_{\mathsf{F}_2}$ whose kth coordinate is $\psi(\sigma_1^k x)$. A shift-commuting map is continuous if and only if[51] it is given by a *sliding block code*. To define this notion, we write $x_{[a,b]}$ for $a \leqslant b$ to denote the block $x_a x_{a+1} \cdots x_b$ of elements of A found in an element x of a shift space between indices a and b. A *block map* is a map Φ defined on the finite collection

$$\{x_{[-m,n]} \mid x \in \mathsf{X}_{\mathsf{F}_1}\}$$

of blocks of length $m+n-1$ seen in points of the shift $\mathsf{X}_{\mathsf{F}_1}$ and taking values in the alphabet A_2 of $\mathsf{X}_{\mathsf{F}_2}$. Such a map defines a sliding block code $\phi : \mathsf{X}_{\mathsf{F}_1} \to \mathsf{X}_{\mathsf{F}_2}$ by

$$(\phi(x))_i = \Phi\left(x_{[i-m,i+n]}\right)$$

for $i \in \mathbb{Z}$. By construction, the map ϕ is shift-commuting and continuous. If ϕ is a homeomorphism then the two shifts are *topologically conjugate*.

The construction in Example A.5 is extremely flexible, and includes a wide array of different phenomena[52]. Some special cases will illustrate this.

(1) If F is empty, then $\mathsf{X}_{\mathsf{F}} = A^{\mathbb{Z}}$, the space of bi-infinite sequences over the finite alphabet A. The resulting dynamical system is called the *full shift* on $|A|$ symbols.

(2) Let $A = \{0, 1\}$ and $\mathsf{F} = \{10^{2k-1}1 \mid k \geqslant 1\}$, where we write 10^n1 to denote the concatenation

$$1\underbrace{0\cdots0}_{n}1.$$

The resulting shift space consists of all bi-infinite binary sequences with the property that there is an even number of 0s between any two 1s, and is called the *even shift*.

(3) Let $A = \{0, 1, 2\}$ and let X denote the set of all bi-infinite sequences over A with the property that a finite block of the form 01^m2^n0 appears in the sequence only if $m = n$. This shift is the *context-free shift*.

A particularly important class of shift dynamical systems are those that can be described by forbidding a finite collection of words.

Example A.6. A shift X with the property that there is a finite set of forbidden words F with $X = \mathsf{X}_{\mathsf{F}}$ is called a *shift of finite type*[53].

(1) The full shift is of finite type, corresponding to an empty set of forbidden words.
(2) Let $\mathsf{F} = 11$ and $A = \{0, 1\}$. Then the resulting shift X_F is of finite type. It consists of the bi-infinite binary sequences with no adjacent 1s. This shift is called the *golden mean* shift.

We will discuss shifts of finite type in more detail in Section A.4.

The final example in this section is a general algebraic construction that subsumes many different types of dynamical system.

Example A.7. Let G be a locally compact group containing a lattice Γ (that is, a discrete subgroup with the property that the quotient space $X = \Gamma\backslash G$ has finite volume with respect to the Haar measure m_G; equivalently with the property that Γ has a fundamental domain in G of finite Haar measure). The Haar measure m_G on G induces a measure m_X on X. For any $g \in G$, the map $T_g : \Gamma h \longmapsto \Gamma h g^{-1}$ is a transformation preserving the measure m_X.

Examples of this sort form part of the study of *homogeneous dynamics*, and a great diversity of phenomena are included.

(1) If $G = \mathbb{R}^d$ and $\Gamma = \mathbb{Z}^d$ then the quotient space $X = \mathbb{T}^d$ is a group, the d-dimensional torus. The map obtained is then a rotation by $g \in \mathbb{R}^d$.
(2) Let $G = \mathrm{PSL}_2(\mathbb{R})$ and $\Gamma = \mathrm{PSL}_2(\mathbb{Z})$. The resulting space $X = \Gamma\backslash G$ can be identified with the unit tangent bundle of a surface of negative curvature (see [53, Ch. 9] for the details), and the time-1 map of the geodesic flow corresponds to the map T_g where
$$g = \begin{pmatrix} \mathrm{e}^{\frac{1}{2}} & 0 \\ 0 & \mathrm{e}^{-\frac{1}{2}} \end{pmatrix},$$
and the time-1 map of the horocycle flow corresponds to the matrix
$$g = \begin{pmatrix} 1 & 1 \\ 0 & 1 \end{pmatrix}.$$

In this volume we will often assume that $\Gamma\backslash G$ is compact, as this simplifies many discussions. Unfortunately, this rules out many natural quotients, $\mathrm{SL}_d(\mathbb{Z})\backslash \mathrm{SL}_d(\mathbb{R})$ in particular. For a treatment of these cases we refer to the lecture notes of Einsiedler and Lindenstrauss [48] and the planned volume [51].

Exercises for Section A.3

Exercise A.3.1. Prove the characterizations of ergodicity for compact group rotations in Example A.3 and for compact group endomorphisms in A.4.

Exercise A.3.2. Describe a set of words F for which X_F is the context-free shift.

A.4 Shifts of Finite Type

There are many convenient ways to describe a shift of finite type, two of which will be particularly useful.

(1) A finite directed graph is a tuple $G = (V_G, E_G)$, where V_G is a finite set whose elements are called vertices, and E_G is a finite set of edges. Each edge $e \in E_G$ has an initial or starting point $i(e) \in V_G$ and a final or terminal point $t(e) \in V_G$. There may be many edges with the same initial and terminal vertex, and an edge may have the same initial and terminal vertex. The structure of such a graph may be recorded via its *adjacency matrix* A_G. This is a $|V_G| \times |V_G|$ matrix whose (i, j)th entry is the number of edges with initial vertex i and terminal vertex j. Given such a graph, we may associate to it the *edge shift*

$$\mathsf{X}_G^{(e)} = \{x = (x_n) \in E_G^{\mathbb{Z}} \mid t(x_n) = i(x_{n+1}) \text{ for all } n \in \mathbb{Z}\}.$$

This is the set of all sequences obtained as the sequence of edges traversed on an infinite path on G.

(2) Assume now that G is a graph with the property that there is no more than one edge between any two vertices (equivalently, with the property that the adjacency matrix A_G has entries in $\{0, 1\}$ only). Then there is an associated *vertex shift* defined by

$$\mathsf{X}_G^{(v)} = \{x = (x_n) \in V_G^{\mathbb{Z}} \mid (A_G)_{x_n x_{n+1}} = 1 \text{ for all } n \in \mathbb{Z}\}.$$

This is the set of sequences obtained by recording the vertices visited while following an infinite path on G.

Thus the graph G, with adjacency matrix

$$A_G = \begin{pmatrix} 1 & 1 \\ 1 & 0 \end{pmatrix},$$

illustrated in Figure A.1 can be used to describe two different shifts. The edge shift $\mathsf{X}_G^{(e)}$ consists of the set of all bi-infinite sequences over the alphabet $\{a, b, c\}$ with the property that the symbol a is always followed by the symbol b, while b can be followed by either a or by c, and c can be followed by a or c. On the other hand, the vertex shift $\mathsf{X}_G^{(v)}$ is the golden mean shift, consisting of all binary sequences with no adjacent 1s. The two shifts are topologically conjugate and both describe simply the set of bi-infinite paths on the graph.

A.4.1 Reduction to Vertex Shifts

In order to relate general shifts of finite type to shifts defined using graphs, we need to assemble some material on different ways to describe the same

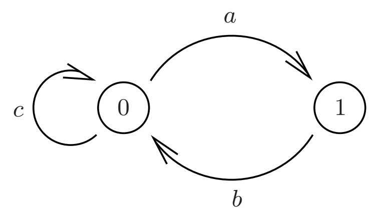

Fig. A.1. The directed graph G.

shift of finite type (up to topological conjugacy). This material is only needed to show that every shift of finite type is topologically conjugate to a vertex shift, and it will not be needed elsewhere. We refer once more to Lind and Marcus [119] for the details.

We say that a shift of finite type is *m-step* if it has a description as X_F where F consists of words of length $(m+1)$. Thus the full shift is 0-step (by convention; here a symbol in F has the same effect as removing that symbol from the alphabet) and the golden mean shift is 1-step, for example.

For any shift $X \subseteq A^{\mathbb{Z}}$ and $N \geqslant 1$, we define an associated shift $X^{[N]}$, the *Nth higher block shift* of X, as follows. Define a new alphabet $A^{[N]}$ to be the collection of all possible words of length N appearing in sequences in X. Define a map ϕ_N, the *Nth higher block code*, from X into $(A^{[N]})^{\mathbb{Z}}$ by sending a sequence x to the sequence whose nth coordinate is the block

$$x_n x_{n+1} \cdots x_{n+N-1} \in A^{[N]}.$$

The Nth higher block shift $X^{[N]}$ of X is then defined to be the image of ϕ_N. We take from Lind and Marcus [119, Sect. 1.4, Prop. 2.3.9] the following results, which we only need as a rationale to restrict to vertex shifts.

- The map ϕ_N is a topological conjugacy between X and $X^{[N]}$.
- The image of ϕ_N is also a shift, and if X is a shift of finite type then $X^{[N]}$ is also.
- Any 1-step shift may be described as a vertex shift, and any vertex shift is a 1-step shift.
- Any edge shift may be described as a vertex shift.
- If X is an m-step shift, then $X^{[m]}$ is a 1-step shift (and hence a vertex shift).
- If X is an m-step shift, then there is a finite directed graph G for which $X^{[m]} = X_G^{(v)}$ and $X^{[m+1]} = X_G^{(e)}$.

From now on we will always work with vertex shifts, which we now describe in more detail.

A.4.2 Markov Measures on Shifts of Finite Type

By the discussion in Section A.4.1 we can reduce the discussion of general shifts of finite type (up to topological conjugacy) to the study of vertex shifts. So let

$$G = (V_G, E_G)$$

be a finite directed graph with the property that the adjacency matrix A_G has entries in $\{0,1\}$ only. For $m \times n$ real matrices $B = (b_{i,j})$ and $C = (c_{i,j})$ we write $B \leqslant C$ (or $B < C$) if $b_{i,j} \leqslant c_{i,j}$ (resp. $b_{i,j} < c_{i,j}$) for $1 \leqslant i \leqslant m$ and $1 \leqslant j \leqslant n$.

We fix a probability vector† $\mathbf{p} = (p_1, \dots, p_v)$ with $v = |V_G|$ coordinates (which we may consider as a probability measure on V_G) and a probability matrix‡ $P \leqslant A_G$. We assume in addition that $\mathbf{p}P = \mathbf{p}$, that is $\mathbf{p}$ is a left eigenvector for P for the eigenvalue 1. This data determines a *Markov measure* $\mu_{\mathbf{p},P}$ on $V_G^{\mathbb{Z}}$ defined by the property

$$\mu_{\mathbf{p},P}\left([a_0 \dots a_s]_n^{n+s}\right) = p_{a_0} P_{a_0,a_1} P_{a_1,a_2} \cdots P_{a_{s-1},a_s} \tag{A.3}$$

where $[a_0 \dots a_s]_n^{n+s}$ is the cylinder set

$$[a_0 \dots a_s]_n^{n+s} = \{x \in \mathsf{X}_G^{(v)} \mid x_n = a_0, x_{n+1} = a_1, \dots, x_{n+s} = a_s\}.$$

As P is a probability matrix, we have

$$\begin{aligned}\sum_{a_s} \mu_{\mathbf{p},P}\left([a_0 \dots a_s]_n^{n+s}\right) &= p_{a_0} P_{a_0,a_1} \cdots P_{a_{s-2},a_{s-1}} \Bigg(\underbrace{\sum_{a_s} P_{a_{s-1},a_s}}_{=1} \Bigg) \\ &= \mu_{\mathbf{p},P}\left([a_0 \dots a_{s-1}]_n^{n+s-1}\right),\end{aligned}$$

which is consistent with the fact that

$$[a_0 \dots a_{s-1}]_n^{n+s-1} = \bigsqcup_{a_s} [a_0 \dots a_s]_n^{n+s}.$$

Since we assume that $\mathbf{p}$ is a left eigenvector of P with eigenvalue 1, that is

$$\sum_{i=1}^{n} p_i P_{i,j} = p_j$$

for $1 \leqslant j \leqslant n$, it follows that

† A vector of non-negative numbers with sum equal to 1.

‡ A square matrix $P = (P_{i,j})$ with non-negative entries and with $\sum_{j=1}^{v} P_{i,j} = 1$ for all i.

$$\sum_{a_0} \mu_{\mathbf{p},P}\left([a_0 \dots a_s]_n^{n+s}\right) = \Bigg(\underbrace{\sum_{a_0} p_{a_0} P_{a_0,a_1}}_{=p_{a_1}}\Bigg) P_{a_1,a_2} \cdots P_{a_{s-1},a_s}$$
$$= \mu_{\mathbf{p},P}\left([a_1 \dots a_s]_{n+1}^{n+s}\right),$$

which is consistent with the fact that

$$[a_1 \dots a_s]_{n+1}^{n+s} = \bigsqcup_{a_0} [a_0 \dots a_s]_n^{n+s}.$$

Using this argument and the Kolmogorov extension theorem [53, Th. A.11], we obtain a well-defined measure on $\mathsf{X}_G^{(v)}$. Since the right-hand side of (A.3) does not depend on n, the measure $\mu_{\mathbf{p},P}$ is invariant under the shift map, resulting in a measure-preserving system $(\mathsf{X}_G^{(v)}, \mu_{\mathbf{p},P}, \sigma)$.

The condition $P \leqslant A_G$ is necessary – it simply means that the measure only gives positive weight to transitions between vertices connected by edges in the graph G.

A.4.3 Ergodicity and Mixing

Markov measures give a class of measure-preserving systems, and we briefly review their basic ergodic properties here. We follow Walters [199, Ch. 1] closely for the material in this section. Write $(a_{i,j}^{(n)})$ for the matrix A^n. A non-negative matrix A is *irreducible* if for each i, j there is some $n = n(i,j) > 0$ with $a_{i,j}^{(n)} = (A^n)_{i,j} > 0$.

Before considering ergodicity for Markov shifts, we assemble some notation and a convenient characterization of ergodicity. Recall that a measure-preserving system $(X, \mathscr{B}, \mu, \sigma)$ is ergodic if and only if

$$\frac{1}{N} \sum_{n=0}^{N-1} \mu\left(\sigma^{-n} A \cap B\right) \longrightarrow \mu(A)\mu(B) \tag{A.4}$$

as $N \to \infty$ for any sets A and B in a semi-algebra that generates $\mathscr{B}$ (see [53, Sect. 2.7]).

Let $(\mathsf{X}_G^{(v)}, \mu_{\mathbf{p},P}, \sigma)$ be as in Section A.4.2 and assume that $\mathbf{p} > 0$. Let us write $f_j = \mathbb{1}_{{}_0[j]_0}$ for the indicator function of the cylinder set

$${}_0[j]_0 = \{x \in X \mid x_0 = j\}.$$

By the pointwise ergodic theorem,

$$\frac{1}{N} \sum_{n=0}^{N-1} f_j\left(\sigma^n x\right) \to f_j^*(x)$$

almost everywhere for some σ-invariant L^∞_μ function f_j^*. Multiplying by f_i, integrating with respect to μ, and applying the dominated convergence theorem shows that

$$\frac{1}{N}\sum_{n=0}^{N-1} p_i p_{ij}^{(n)} \to \int_X f_i(x) f_j^*(x) d\mu.$$

It follows that

$$\frac{1}{N}\sum_{n=0}^{N-1} P^n \longrightarrow Q = (q_{ij})$$

as $N \to \infty$, where $q_{ij} = p_i^{-1} \displaystyle\int_X f_i(x) f_j^*(x) d\mu$. It is straightforward to check that Q is a stochastic matrix with $QP = PQ = Q$, $Q^2 = Q$, and that any eigenvector of P with eigenvalue 1 is also an eigenvector of Q.

Theorem A.8. *For the measure-preserving system* $(X, \mu, \sigma) = (\mathsf{X}_G^{(v)}, \mu_{\mathbf{p},P}, \sigma)$ *with* $\mathbf{p} > 0$ *the following properties are equivalent.*

(1) *the system is ergodic;*
(2) *every row of the matrix* Q *defined above is identical;*
(3) $q_{i,j} > 0$ *for every* i, j*;*
(4) *the matrix* P *is irreducible; and*
(5) 1 *is a simple eigenvalue of* P *(in the sense that* P *has a one-dimensional eigenspace for the eigenvalue* 1*).*

Proof. Assume that the system is ergodic as in (1), so that

$$\lim_{N\to\infty} \frac{1}{N}\sum_{n=0}^{N-1} \mu\left([i]_0 \cap [j]_n\right) = p_i q_{i,j}$$

by the argument above, and by the characterization of ergodicity in (A.4) we have $p_i q_{i,j} = p_i p_j$, showing (2).

Assuming that the rows of Q are identical as in (2), $\mathbf{p}Q = \mathbf{p}$ implies that $q_{i,j} = p_j > 0$, and hence (3).

Assume now $Q > 0$ as in (3). By definition of Q, we have for any given i, j,

$$\lim_{N\to\infty} \frac{1}{N}\sum_{n=0}^{N-1} p_{i,j}^{(n)} = q_{i,j} > 0,$$

and therefore (3) implies that there must be some n with $p_{i,j}^{(n)} > 0$, showing that P is irreducible as in (4).

Assume now that P is irreducible. Fix some i and j and recall that

$$Q = QP^n$$

for any n. Therefore $q_{i,j} \geqslant q_{i,\ell} p^{(n)}_{\ell,j}$ for all ℓ. As Q is stochastic there exists some ℓ with $q_{i,\ell} > 0$. As P is assumed to be irreducible there exists some n such that $p^{(n)}_{\ell,j} > 0$, which shows that $q_{i,j} > 0$ as required in (3).

Assume $q_{i,j} > 0$ for all i, j as in (3), fix j and let $q_j = \max_i q_{i,j}$. If $q_{i',j} < q_j$ for some i' then, since $Q^2 = Q$,

$$q_{\ell,j} = \sum_i q_{\ell,i} q_{i,j} < \sum_i q_{\ell,i} q_j = q_j$$

for any ℓ, contradicting the choice of q_j. Thus (2) holds.

Finally, assume that (2) holds, and recall that $\mathbf{p}Q = \mathbf{p}$ implies that $q_{i,j} = p_j$ for all i, j. By (A.4) it is enough to check the convergence on cylinder sets

$$A = [i_0, \dots, i_r]_a^{a+r}$$

and

$$B = [j_0, \dots, j_s]_b^{b+s}.$$

For sufficiently large n we have

$$\mu\left(\sigma^{-n} A \cap B\right) = p_{j_0} \prod_{k=0}^{s} p_{j_k, j_{k+1}} p^{(n+a-b-s)}_{j_s, i_0} \prod_{k=0}^{r-1} p_{i_k, i_{k+1}} \tag{A.5}$$

by (A.3). Therefore by definition of Q we get

$$\lim_{N\to\infty} \frac{1}{N} \sum_{n=0}^{N-1} \mu\left(\sigma^{-n} A \cap B\right) = \mu(A)\mu(B).$$

To see that (2) and (5) are equivalent, we argue as follows. Since $Q = QP$ each row of Q is a left eigenvector of P, and so if 1 is a simple eigenvalue of P as in (5) then (2) holds. On the other hand (2) gives that Q has rank one, and by definition of Q every left eigenvector of P for eigenvalue 1 is also an eigenvector for Q for eigenvalue 1. Therefore, P must have 1 as a simple eigenvalue. □

Mixing is a stronger property for a Markov shift. A non-negative matrix A is *primitive* if there is some $m \geqslant 1$ with $a^{(m)}_{i,j} > 0$ for all i, j.

Example A.9. Let $A = \begin{pmatrix} 2 & 1 \\ 0 & 1 \end{pmatrix}$, $B = \begin{pmatrix} 0 & 2 \\ 1 & 0 \end{pmatrix}$ and $C = \begin{pmatrix} 0 & 2 \\ 1 & 1 \end{pmatrix}$. Then A is non-negative but not irreducible, B is irreducible but not primitive, and C is primitive.

Theorem A.10. *For the measure-preserving system* $(X, \mu, \sigma) = (\mathsf{X}^{(v)}_G, \mu_{\mathbf{p},P}, \sigma)$ *with* $\mathbf{p} > 0$ *the following properties are equivalent.*

(1) *the system is mixing;*

(2) $p_{i,j}^{(n)} \to p_j$ *for all* i,j *as* $n \to \infty$*; and*
(3) *the matrix* P *is primitive.*

We will use the following lemma for the proof.

Lemma A.11. *Let* P *be a primitive stochastic matrix, let* $\mathbf{p}$ *be a stochastic vector with* $\mathbf{p}P = \mathbf{p}$*, and let*

$$L = \left\{ \mathbf{x} = (x_1, \dots, x_v) \mid \sum_{i=1}^{v} x_i = 0 \right\}.$$

Then the eigenvalues of the linear map $P|_L : L \to L$ *sending* $\mathbf{x}$ *to* $\mathbf{x}P$ *all have absolute value less than one.*

Proof. First notice that $\mathbf{x} \in L$ implies that

$$\sum_i (\mathbf{x}P)_i = \sum_{i,j} x_j p_{ji} = \sum_j x_j = 0$$

as P is stochastic. Hence L is invariant under P.

Assume that $P^m > 0$ for some $m \geqslant 1$. Then the continuous map $\mathbf{x} \mapsto \mathbf{x}P$ sends the set

$$\Delta = \{ \mathbf{x} = (x_1, \dots, x_v) \mid \mathbf{x} \geqslant 0, \sum_{i=1}^{v} x_i = 1 \}$$

to itself[†] and by assumption fixes a probability vector $\mathbf{p} = \mathbf{p}P^m > 0$. Since P^m is strictly positive,

$$\Delta P^k \subseteq \Delta P^m \subseteq \Delta^o \tag{A.6}$$

for all $k \geqslant m$.

Now argue as follows.

- Since $\|\mathbf{x}P\|_1 \leqslant |\mathbf{x}|_1$, no eigenvalue can be greater than 1 in modulus.
- So suppose some $\omega \in \mathbb{C}$ with $|\omega| = 1$ is an eigenvalue for $P|_L$, which shows that there is a sequence $k_n \to \infty$ with $\omega^{k_n} \to 1$ as $n \to \infty$.
- This shows that there exists some nonzero $\mathbf{v} \in L$ with $\mathbf{v}P^{k_n} \to \mathbf{v}$. We may normalize the length of $\mathbf{v}$ so that $\mathbf{p} + \mathbf{v} \in \partial\Delta$. However, this contradicts (A.6).

□

Proof of Theorem A.10. We will make use of characterizations of the various mixing properties from [53, Ch. 2]. Assume (1), so

$$p_i^{-1}\Big(\mu([i]_0 \cap \sigma^{-n}[j]_0) - \mu([i]_0)\mu([j]_0)\Big) = p_{i,j}^{(n)} - p_j \to 0$$

[†] A consequence is that the map has a fixed point by the Brouwer fixed point theorem, and this fixed point is the probability vector $\mathbf{p}$.

as $n \to \infty$ for all i, j, showing (2), which clearly implies (3).

Assume (3) so that P is a primitive stochastic matrix with stochastic eigenvector $\mathbf{p}$. For fixed i, notice that the ith basis vector e_i lies in Δ, so that

$$p_{i,j}^{(n)} = (e_i P^n)_j$$

and we may write

$$e_i = \mathbf{p} + \underbrace{(e_i - \mathbf{p})}_{=\mathbf{s} \in L}.$$

It follows by Lemma A.11 that $\mathbf{s}P^n \to 0$ as $n \to \infty$, giving (2). Mixing for cylinder sets now follows by (A.5), and this shows the Markov shift is mixing because we can approximate any measurable set by a finite union of cylinder sets. □

Exercises for Section A.4

Exercise A.4.1. Let $X = \mathsf{X}_G^{(e)}$ be the edge graph associated to the directed graph G, and let σ be the shift map on X.
(a) Show that

$$|\{x \in X \mid \sigma^n x = x\}| = \operatorname{tr}(A_G^n).$$

(b) Deduce that the dynamical zeta function of σ, defined by

$$\zeta_\sigma(z) = \exp \sum_{n=1}^{\infty} |\{x \in X \mid \sigma^n x = x\}| \frac{z^n}{n}$$

for $s \in \mathbb{C}$, has a positive radius of convergence and defines a rational function where it converges.
(c) Find the radius of convergence of the dynamical zeta function in terms of the eigenvalues of the adjacency matrix A_G.

Exercise A.4.2. Find a graph G with the property that $\mathsf{X}_G^{(v)}$ is topologically conjugate to the shift of finite type $\mathsf{X}_{\{000,111\}}$.

Exercise A.4.3. Show that if a Markov shift $(\mathsf{X}_G^{(v)}, \mu_{\mathbf{p},P}, \sigma)$ is weak-mixing, then it is mixing.

A.5 Invertible Extensions

It is often useful to assume that a measure-preserving system is invertible. In this short section we describe how to build an invertible system from any measure-preserving system on a Borel probability space[†]. This construction appears as part of the foundational work of Rokhlin [174, Sect. 3]. Let $\mathsf{X} = (X, \mathscr{B}_X, \mu, T)$ be any measure-preserving system. Then the system $\widetilde{\mathsf{X}} = (\widetilde{X}, \widetilde{\mathscr{B}}, \widetilde{\mu}, \widetilde{T})$ defined by

- $\widetilde{X} = \{x \in X^{\mathbb{Z}} \mid x_{k+1} = T(x_k) \text{ for all } k \in \mathbb{Z}\}$;
- $(\widetilde{T}(x))_k = x_{k+1}$ for all $k \in \mathbb{Z}$ and $x \in \widetilde{X}$;
- $\widetilde{\mu}\left(\{x \in \widetilde{X} \mid x_0 \in A\}\right) = \mu(A)$ for any $A \in \mathscr{B}_X$, and $\widetilde{\mu}$ is invariant under $\widetilde{T}$;
- $\widetilde{\mathscr{B}}$ is the smallest $\widetilde{T}$-invariant σ-algebra for which the map $\pi : x \mapsto x_0$ from $\widetilde{X}$ to X is measurable;

is an invertible measure-preserving system, and the map $\pi : x \mapsto x_0$ is a factor map. The system $\widetilde{\mathsf{X}}$ is called the *invertible extension* of X.

The invertible extension $\widetilde{\mathsf{X}}$ has the following universal property. For any extension

$$\phi : (Y, \mathscr{B}_Y, \nu, S) \to (X, \mathscr{B}_X, \mu, T)$$

for which S is invertible, there exists a unique map

$$\widetilde{\phi} : (Y, \mathscr{B}_Y, \nu, S) \to (\widetilde{X}, \widetilde{B}, \widetilde{\mu}, \widetilde{T})$$

for which $\phi = \pi \circ \widetilde{\phi}$.

Outline of Proof. Let $X \subseteq \bar{X}$ be a Borel subset of the compact metric space $\bar{X}$. As the σ-algebra $\mathscr{B}_{\bar{X}}$ is countably generated one can check that $\{x \in \bar{X}^{\mathbb{Z}} \mid x_0 \in X \text{ and } x_1 = T(x_0)\}$ is a Borel subset of the compact metric space $\bar{X}^{\mathbb{Z}}$. This implies also that $\widetilde{X} \subseteq \overline{X}^{\mathbb{Z}}$ is a Borel subset. Also show that $x \in X \mapsto (x, T(x), T^2(x), \dots) \in \bar{X}^{\mathbb{N}_0}$ is measurable.

Using μ we now define a measure μ_0 on $\bar{X}^{\mathbb{N}_0}$ by the formula

$$\mu_0(B) = \mu(\{x \in X : (x, T(x), T^2(x), \dots) \in B\})$$

for every Borel subset $B \subseteq \bar{X}^{\mathbb{N}_0}$. Lift μ_0 to the measure $\widetilde{\mu}_0 = \delta_z \times \mu_0$, where δ_z denotes the Dirac measure at some point $z \in \bar{X}^{-\mathbb{N}}$. We write $\sigma : \bar{X}^{\mathbb{Z}} \to \bar{X}^{\mathbb{Z}}$ for the shift map defined by $\sigma(x)_n = x_{n+1}$. Define $\widetilde{\mu}_n = \sigma_*^n \widetilde{\mu}_0$ for $n \geqslant 0$ and show (e.g. using Tychonoff-Alaoglu) that $\widetilde{\mu}_n$ converges in the weak* topology to a probability measure $\widetilde{\mu}$ on $\bar{X}^{\mathbb{Z}}$ that satisfies $\widetilde{\mu}(\widetilde{X}) = 1$. Now check that $\widetilde{\mu}$ is invariant under the shift map and conclude the remaining properties. □

For measure-preserving systems with additional structure, it is often possible to preserve some features in the invertible extension. This sometimes involves *ad hoc* constructions, but for continuous maps on compact metric spaces the following simple construction suffices. Assume that $T : X \to X$ is a continuous map on a compact metric space (X, d). We define the invertible extension $\widetilde{T} : \widetilde{X} \to \widetilde{X}$ as follows. Let

[†] A *Borel probability space* is a Borel subset X of a compact metric space $\overline{X}$, with a probability measure μ defined on the restriction of the Borel σ-algebra $\mathscr{B}$ to X.

- $\widetilde{X} = \{x \in X^{\mathbb{Z}} \mid x_{k+1} = Tx_k \text{ for all } k \in \mathbb{Z}\}$;
- $(\widetilde{T}x)_k = x_{k+1}$ for all $k \in \mathbb{Z}$ and $x \in \widetilde{X}$;

with metric $\widetilde{\mathsf{d}}(x,y) = \sum_{k\in\mathbb{Z}} 2^{-|k|}\mathsf{d}(x_k,y_k)$. Write $\pi : \widetilde{X} \to X$ for the map sending x to x_0. We can then check the following properties:

(1) $\widetilde{T}$ is a homeomorphism of the compact metric space $\widetilde{X}$, and $\pi : \widetilde{X} \to X$ is a topological factor map.
(2) If (Y,S) is any homeomorphism of a compact metric space with the property that there is a topological factor map $(Y,S) \to (X,T)$, then $(\widetilde{X},\widetilde{T})$ is a topological factor of (Y,S).
(3) $\pi_*\mathscr{M}^{\widetilde{T}}(\widetilde{X}) = \mathscr{M}^T(X)$.
(4) $\pi_*\mathscr{E}^{\widetilde{T}}(\widetilde{X}) = \mathscr{E}^T(X)$.

A.6 The Rokhlin–Halmos Lemma

The Kakutani–Rokhlin lemma [172] (see [53, Lem. 2.45]) for ergodic measure-preserving systems was extended to the setting of aperiodic transformations by Halmos [79].

Definition A.12. *A measure-preserving system* $(X,\mathscr{B},\mu,T)$ *is* aperiodic *if*

$$\mu(\{x \in X \mid T^n x = x\}) = 0$$

for all $n \in \mathbb{N}$.

Theorem A.13 (Rokhlin–Halmos lemma). *Let* $(X,\mathscr{B},\mu,T)$ *be an invertible aperiodic measure-preserving system on a Borel probability space. Then for any* $n \geqslant 1$ *and* $\varepsilon > 0$ *there is a set* $B \in \mathscr{B}$ *for which*

$$B, TB, \dots, T^{n-1}B$$

are disjoint, and

$$\mu\left(B \sqcup TB \sqcup \cdots \sqcup T^{n-1}B\right) > 1 - \varepsilon.$$

The proof proceeds along more or less the same lines as the proof of the Kakutani–Rokhlin lemma for ergodic systems, analysing the properties of *Rokhlin towers* (or *Kakutani skyscrapers*), being sets of the form

$$A \cup TA \cup \cdots \cup T^{n-1}A$$

for measurable sets A chosen to make the union a disjoint one. The absent hypothesis of ergodicity is roughly speaking replaced with a notion of maximality among measurable sets with a property of this sort (Lemma A.15).

Lemma A.14. *If $(X, \mathscr{B}, \mu, T)$ is an invertible aperiodic measure-preserving system on a Borel probability space, and $A \in \mathscr{B}$ has $\mu(A) > 0$, then for any $n \geqslant 1$ there is a measurable set $F \subseteq A$ of positive measure for which*

$$F, TF, \dots, T^{n-1}F \tag{A.7}$$

are disjoint.

Proof. The statement holds trivially for $n = 1$, since in this case we may take $F = A$. Assume for the purposes of an induction argument that the statement holds for n and let $F \subseteq A$ be the corresponding subset so that $F, \dots, T^{n-1}A$ are all disjoint. We claim that there is a measurable set $B \subseteq F$ such that

$$\mu(B \triangle T^n B) > 0.$$

If not, then for any measurable set $B \subseteq F$ we have $\mu(B \triangle T^n B) = 0$, so T^n coincides modulo μ with the identity on F, contradicting the hypothesis of aperiodicity (see Exercise A.6.2). Define $F' = B \smallsetminus T^n B$. Then

$$F', TF', \dots, T^n F'$$

are all disjoint, and

$$\mu(F') = \mu(B) - \mu(G \cap T^n B) > 0.$$

This shows the statement of the lemma for $n + 1$, and hence the lemma by induction. □

Lemma A.15. *Let $(X, \mathscr{B}, \mu, T)$ be an invertible aperiodic measure-preserving system on a Borel probability space. Then for any $n \geqslant 1$ there exists a measurable set F of positive measure for which*

$$F, TF, \dots, T^{n-1}F$$

are disjoint, and that is maximal with respect to that property[†].

Proof. We wish to find F as a countable union $F = \bigcup_{k \geqslant 1} F_k$ of a sequence of measurable sets such that $F_k \subseteq F_{k+1}$ and $F_k, TF_k, \dots, T^n F_k$ are disjoint for all $k \geqslant 1$.

For any $B \in \mathscr{B}$ for which $B, TB, \dots, T^{n-1}B$ are disjoint we define

$$s(B) = \sup\bigl\{\mu(G) : G \supseteq B \text{ and } G, TG, \dots, T^{n-1}G \text{ are disjoint}\bigr\} \geqslant \mu(B).$$

Let F_1 be any set with the disjointness property and positive measure as in Lemma A.14.

[†] That is, if $G \supseteq F$ has the same properties, then $\mu(G) = \mu(F)$.

If F_k has already been constructed, we define $s_k = s(F_k)$. If $s_k = \mu(F_k)$ then the lemma has been proven. If not, then there exists some $F_{k+1} \supseteq F_k$ as in the definition of s_k with $\mu(F_{k+1}) > \frac{\mu(F_k)+s_k}{2}$.

Note that $s_{k+1} \leqslant s_k$, which by induction implies $s_k - \frac{1}{2^{k-1}} < \mu(F_k) \leqslant s_k$. This implies for $F = \bigcup_{k\geqslant 1} F_k$ that $F, TF, \ldots, T^{n-1}F$ are disjoint and

$$\mu(F) = \lim_k \mu(F_k) = \lim_k s_k = s(F).$$

□

Proof of Theorem A.13. For a fixed $m \geqslant 1$ let F be a maximal set constructed as in Lemma A.15, and let $r : F \to \mathbb{N}$ denote the first return time to the set F. By the disjointness property we have $r(x) \geqslant m$ for $x \in F$.

We define $F_k^* = \{x \in F \mid r(x) = k\}$ for all $k \geqslant m$ so that

$$F \underset{\mu}{=} \bigsqcup_{k=m}^{\infty} F_k^*$$

by Poincaré recurrence†. Moreover, we claim that

$$X' = \bigsqcup_{k=m}^{\infty} \bigsqcup_{i=0}^{k-1} T^i F_k^* \underset{\mu}{=} X. \tag{A.8}$$

That the union is disjoint follows easily from the construction of F_k^*. To see the last equation in (A.8), we first show that X' is a T-invariant set as follows. If $x \in T^i F_k^*$ for some $k \geqslant m$ and some i with $0 \leqslant i < k-1$, then $Tx \in T^{i+1}F_k^* \subseteq X'$ by definition of X'. On the other hand, if $x \in T^{k-1}F_k^*$ for some $k \geqslant m$ then $Tx \in T^k F_k^* \subseteq F \subseteq X'$ by definition of F_k^*. It follows that $X \smallsetminus X'$ is also T-invariant, and if $X \smallsetminus X'$ has positive measure then we can apply Lemma A.14 to find a measurable set $F' \subseteq X \smallsetminus X'$ with positive measure for which $F', TF', \ldots, T^{m-1}F'$ are disjoint. This means that $F \cup F'$ has strictly larger measure than F, contradicting the maximality property. Thus $\mu(X \smallsetminus X') = 0$ as claimed. Notice that this argument is using the maximality property as a substitute for ergodicity.

Let $n \geqslant 1$ and $\varepsilon > 0$ be arbitrary. We will construct B as in the theorem using the above construction and division with remainder in $\mathbb{N}$. In fact, we choose $m \geqslant 1$ with $\frac{n}{m} < \varepsilon$, and let F be as above. We claim that the set

$$B = \bigcup_{k=m}^{\infty} \bigcup_{0\leqslant i\leqslant \lfloor \frac{k}{n} \rfloor - 1} T^{in} F_k^*$$

has the properties as in the theorem. By construction the sets

† The definition of F implies that it is sufficient to take the union for k running from m to $2m-1$, but we will not need this in the proof.

$$B, TB, \dots, T^{n-1}B$$

are disjoint and

$$\bigcup_{j=0}^{n-1} T^j B = \bigcup_{k=m}^{\infty} \bigcup_{\ell=0}^{n\lfloor \frac{k}{n} \rfloor - 1} T^\ell F_k^*.$$

By (A.8) we get that

$$X \smallsetminus \bigcup_{j=0}^{n-1} T^j B \underset{\mu}{=} \bigcup_{k=m}^{\infty} \bigcup_{\ell=n\lfloor \frac{k}{n} \rfloor}^{k-1} T^\ell F_k^*.$$

Taking the measure we get

$$\mu\left(X \smallsetminus \bigcup_{j=0}^{n-1} T^j B\right) = \sum_{k=m}^{\infty} \left(k - n\left\lfloor \frac{k}{n} \right\rfloor\right) \mu(F_k^*) \leqslant \sum_{k=m}^{\infty} \frac{n}{m} k\mu(F_k^*) < \varepsilon,$$

where we used that for each $k \geqslant m$ the sets $T^i F_k^*$ have the same fixed measure independent of i, that $k \geqslant m$, that

$$\sum_{k \geqslant m} k\mu(F_k^*) = \mu(X') = 1,$$

and our choice of m. This proves the theorem. □

Exercises for Section A.6

Exercise A.6.1. Show that if $(X, \mathscr{B}, \mu, T)$ is aperiodic then μ is non-atomic.

Exercise A.6.2. Let $(X, \mathscr{B}, \mu, T)$ be a measure-preserving system on a Borel probability space and let $F \in \mathscr{B}$ have positive measure. If $\mu(B \triangle T^{-1}B) = 0$ for all measurable $B \subseteq F$, show that

$$\mu\left(\{x \in F \mid Tx \neq x\}\right) = 0.$$

Notes to Appendix A

(47)(Page 246) This is a special case of a result of Boshernitzan [18].

(48)(Page 247) This characterization is due to Ornstein [155], who used it in order to construct a mixing transformation T with no square root (that is, with no measure-preserving system S for which S^2 is measurably isomorphic to T). With no mixing hypothesis at all it is easy to construct such examples, and Chacon [34] had earlier found weak-mixing examples. In contrast a consequence of Ornstein's work on Bernoulli shifts is that any Bernoulli shift does have a square root.

(49)(Page 247) The fact that there is always at least one invariant Borel probability measure is the Kryloff–Bogoliouboff Theorem [111].

(50)(Page 249) Theorems of this sort are of course powerful. In addition to the basic examples (the ergodic theorems themselves, the Kryloff–Bogoliouboff Theorem, Poincaré recurrence, for example) there are much deeper general phenomena like the multiple recurrence theorem of Furstenberg [65] (see [53, Ch. 7]).

(51)(Page 251) This result was shown by Hedlund [82, Th. 3.4], and this essential finiteness property has led to significant interactions between symbolic dynamics and information theory. We refer to the papers of Adler, Coppersmith and Hassner [8] and Marcus [129] for an introduction to this. Despite the finite nature of a topological conjugacy, the problem of determining when two shifts of finite type are topologically conjugate is difficult.

(52)(Page 251) Shift dynamical systems, or symbolic dynamics, is a substantial branch of dynamical systems in its own right. Hadamard [77] studied the geodesic flow on a surface of negative curvature, and showed that the trajectories could be coded (by partitioning the space and recording the motion through the elements of the partition), with the set of symbolic sequences arising being described by forbidding finitely many pairs. Symbolic dynamics also arise naturally in data transmission, recording, and compression. Thus, for example, devising a method to record a binary sequence onto a physical medium with constraints (for example, to avoid long runs of the same symbol) in an optimal fashion can be viewed as the problem of understanding the constraints on certain natural maps between shift spaces. The abstract subject of studying symbolic dynamical systems was given enormous impetus by the work of Morse and Hedlund [138]. For accessible modern treatments, consult the monographs of Lind and Marcus [119] for a broad overview, Kitchens [104] for more on the theory of Markov shifts, and the collection edited by Berthé, Ferenczi, Mauduit and Siegel [61] for the rather different theory of shift dynamical systems of low complexity.

(53)(Page 251) Shifts of finite type have a rich structure, and their classification up to the natural notion of equivalence (topological conjugacy via a shift-commuting homeomorphism) is extremely subtle. The monograph of Lind and Marcus [119] is a good introduction to this topic.